arXiv:1610.01458v1 [cs.DM] 5 Oct 2016

Distributed Searching of Partial Grids*

Dariusz Dereniowski† Dorota Urbańska†

Abstract

We consider the following distributed pursuit-evasion problem. A team of mobile agents called searchers starts at an arbitrary node of an unknown n-node network. Their goal is to execute a search strategy that guarantees capturing a fast and invisible intruder regardless of its movements using as few agents as possible. We restrict our attention to networks that are embedded into partial grids: nodes are placed on the plane at integer coordinates and only nodes at distance one can be adjacent. We give a distributed algorithm for the searchers that allow them to compute a connected and monotone strategy that guarantees searching any unknown partial grid with the use of $O(\sqrt{n})$ searchers. As for a lower bound, not only there exist partial grids that require $\Omega(\sqrt{n})$ searchers, but we prove that for each distributed searching algorithm there is a partial grid that forces the algorithm to use $\Omega(\sqrt{n})$ searchers but $O(\log n)$ searchers are sufficient in the offline scenario. This gives a lower bound of $\Omega(\sqrt{n}/\log n)$ in terms of achievable competitive ratio of any distributed algorithm.

keywords: connected search number, distributed searching, graph searching, partial grid, pursuit-evasion

1 Introduction

A team of mobile autonomous robots wants to search an area with the goal of finding a mobile intruder (or lost entity). The intruder has several properties that dictate how a search should be conducted. First, the intruder is invisible and therefore the robots may conclude its potential locations only from the history of their own moves. Second, it is assumed that the speed of the intruder is unknown and therefore the robots build their search strategy assuming that the intruder is very fast: may traverse arbitrarily long distance between any two actions of a robot. Third, the intruder is very clever, i.e., it will avoid being captured as long as possible; in other words we may imagine that it knows locations of robots and their future movements at any point. This assumption enforces robots to consider the worst case scenario for them since they want to have a search strategy that guarantees interception. The above problem is usually restated in discrete

*Research partially supported by National Science Centre, Poland, grant number 2015/17/B/ST6/01887.

†Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, Gdańsk, Poland

terms, naturally expressing the search game using graph-theoretic notation. Following the widely used terminology, the mobile entities performing the search are called *searchers*.

In this work we focus on the graph-theoretic problem statement, where the searchers operate in a given graph in which they move along edges. Moreover, what greatly influences algorithmic approach is assumption whether the searchers know the graph in advance (offline version of the problem) or whether the graph is unknown and the searchers learn its structure while conducting the search (online or distributed setting). We will shortly review both approaches, giving later a formal statement of the problem we study in this work. In all cases we are interested in minimizing the number of searchers needed to clear the given network.*

Off-line searching. The offline graph searching models are extensively studied and numerous deep results have been obtained, providing insight into not only the problem itself but also enriching the more widely understood graph theory through the connections between graph searching games and many graph parameters, e.g., pathwidth, treewidth, branchwidth, bandwidth, profile, interval thickness, vertex separation number; see e.g. [19] for a survey and further references. The historically first studied graph searching model is called *edge search* [29, 30]. In this problem, the goal is to construct a search strategy that guarantees capturing a fast and invisible fugitive (thus, the strategy must ensure success regardless of the moves performed by the fugitive) in a graph that is given as an input to an algorithm computing a search strategy. A search strategy itself is a sequence of *moves*, where each move is one of the following: (i) placing a searcher on any graph node; (ii) removing a searcher from the node it occupies; (iii) sliding a searcher along an edge in order to clear it. Since we adopt the connected searching problem in our distributed model, we point out to few recent works on the problem [1, 2, 10, 11, 15, 16].

Distributed searching. In the distributed, or on-line, version of the problem it is assumed that the network is unknown in advance to the searchers. In this setting, some assumptions need to be made. First, only *monotone* search strategies are considered. In a monotone search strategy, once an edge has been cleared, it must remain clean till the end of the search; in other words, the subgraph composed of edges that may contain the fugitive may only shrink as the search progresses thus disallowing any recontaminations. This assumption is dictated by an observation that otherwise the searchers may first learn the structure of the network by exploring it (and thus ignoring the possibility of capturing the intruder at this stage) and once the network is known, they can compute a search strategy by using an off-line algorithm and finally execute the strategy. The problem then reduces to exploration and map construction, well studied problems in distributed computing. Another natural assumption is to forbid placing a searcher on a node that has not been visited before.

We consider *connected* search strategies in this work, i.e., strategies that guarantee that at any given time point the subgraph that is clean is connected. Note that this allows us to assume that all searchers start at some node called the *homebase* and only moves of type (iii) are then made (see the definition of edge search above). Indeed, removing a searcher from a node u and placing it on another v (i.e., *jumping*) one may be replaced†

*In this work the terms graph and network are used exchangeably.

†Note that this observation is true as long as the considered optimization criterion is to minimize the number of searchers. For other criteria, like e.g. search time, the exclusion of jumping may be potentially limiting.

by a sequence of sliding moves along a path from u to v consisting of clean edges only (such a path must exists due to connectedness and monotonicity).

1.1 Related work

Off-line problems. One of the central questions raised in the context of various graph searching problems is the monotonicity which can be stated as follows. A searching problem is said to be *monotone* if there exists a monotone search strategy solving the problem. Note that proving monotonicity is a tool that allows to conclude membership in NP for a given problem. It is known that the edge search problem is monotone [25]. On the other hand, the connected search is not monotone [36]. A question related to the latter searching model is: how many extra searchers one needs to ensure connectivity. It turns out that each monotone edge search strategy can be converted (in polynomial time) into a monotone connected one by approximately 'doubling' the number of searchers [12]. Thus, for asymptotic results, like the one in this work, this gives another reason that justifies restricting attention to monotone connected search strategies.

Distributed searching. In most cases, when designing distributed searching algorithms, the monotonicity requirement is adopted. (See [4] for an example how an optimal connected search strategy can be constructed in a distributed fashion when recontamination is allowed.) During construction of a monotone strategy in a distributed way, there is naturally some 'cost' involved in terms of increased number of searchers required for guarding — this cost measured as the ratio of number of searchers that each distributed algorithm needs to use for some n-node graph and its monotone connected (off-line) search number is know to be $\Omega(n/\log n)$ [23]. In the realm of distributed algorithms, natural questions arise with respect to the amount (and type) of additional information regarding the underlying network given a priori to an algorithm. In [28] is was proved that $O(n \log n)$ bits of *advice* are sufficient to construct an optimal connected monotone search strategy (the concept of such quantitative approach to advice analysis was introduced in [20]). An example of an algorithmic approaches in a very weak computational model see e.g. [3, 9].

Grid networks were studied in [8] where the searching model used the concept of temporal immunity: a node after cleaning remains protected (even if unguarded) against recontamination for a certain amount of time. For other searching works involving immunity see e.g. [17, 18]. For other distributed searching models and algorithms for specific network topologies see [6, 16, 21, 27]. We provide in Section 7 a brief discussion of a potential applicability of our result in 'continuous' environments, like polygons.

Applications in robotics. We note that our results may be of particular interest not only providing theoretical insight into searching dynamics in distributed agent computations, but may also find applications in the field of robotics. Most investigations oriented towards algorithms that can be applied on physical devices need to deal with the problem of modeling of real world. This can be done either by discretizing it (usually via graph theory) or by building algorithms that work in continuous search space and need to address the geometric issues that emerge. (In Section 7.1, we add a brief discussion on this subject from the point of view of our results.) Having in mind vast literature on the subject we point out interested reader by providing few references to recent works in this field [7, 13, 22, 24, 31, 32, 33, 34, 35].

1.2 Outline of this work

The next section introduces the notation used in this work and provides problem statement. It is subdivided so that Section 2.1 defines the graph searching problem we study while Section 2.2 introduces the terminology related to the partial grid networks we consider in this paper. Section 3 provides a construction of a class of n-node networks such that each distributed algorithm uses $\Theta(\sqrt{n})$ searchers which turns out to be $\Omega(\sqrt{n}/\log n)$ times more that an optimal off-line algorithm would use (recall that by off-line algorithm we refer to the case when the entire network is given as an input and hence is known in advance to the algorithm). This serves as a lower bound in our analysis.

Section 4 describes a distributed algorithm that performs a guaranteed search in partial grids and it is assumed that the algorithm is given an upper bound n on the size of the network. We point out that this algorithm uses a distributed procedure from [5] as a subroutine that is called many times to clear selected parts of a grid, and it can be seen as a generalization from a 'linear' graph structure studied in [5] to a 2-dimensional structure discussed in this work. Also, although both algorithms are conducted via some greedy rules which dictate how a search should 'expand' to unknown parts of the graph, the analysis of our algorithm is different from the one in [5].

Then, in Section 5 we prove the correctness of the algorithm and provide an upper bound on its performance: it is using $O(\sqrt{n})$ searchers for any partial grid network. In Section 6 we consider a modified version of the algorithm, which receives no information on the underlying graph in advance, and we prove that the algorithm also uses $O(\sqrt{n})$ searchers. This result, stated in Theorem 6.1, is the main contribution of this work. We finish with conclusions in Section 7, giving few remarks on how our work relates to searching two-dimensional environments, like polygons with holes. As there are many open problems and research directions related to the subject, we list some of them also in Section 7.

2 Definitions and terminology

In this section we will present the notation we use. We consider only simple undirected connected graphs. Given any graph $G = (V, E)$ and $X \subseteq V$, $G[X]$ is the subgraph of G *induced by* X: its node set is X and consists of all edges $\{u, v\}$ of G having both endpoints in X.

2.1 Problem statement

A *connected k-search strategy* $\mathcal{S}$ for a network G is defined as follows. Initially, k searchers are placed on a node h of G, called the *homebase*. (We also say that $\mathcal{S}$ *starts at* h.) Then, $\mathcal{S}$ is a sequence of moves, where each *move* consists of selecting one searcher present at some node u and sliding the searcher along an edge $\{u, v\}$. (Thus, the searcher moves from its current location to one of the neighbors.) Initially, all edges are *contaminated*.

After each move of sliding a searcher along an edge $\{u, v\}$ it is declared to be *clean*. It becomes contaminated again (*recontaminated*) if at any time during execution of the strategy $\mathcal{S}$ at least one of its endpoints is not occupied by a searcher and is incident to a contaminated edge. If no recontamination happens in $\mathcal{S}$, then $\mathcal{S}$ is called *monotone*. Regardless of whether the strategy is monotone or not, we require that the subgraph

consisting of all clean edges is connected after each move of the search strategy. Finally, we require that after the last move of $\mathcal{S}$ all edges are clean.

The minimum k such that there exist a node h and a (monotone) connected k-search strategy that starts at h is called the (monotone) *connected search number* of G and denoted by ($\texttt{mcs}(G)$, respectively) $\texttt{cs}(G)$.

Having defined a search strategy, we now state the distributed model we use. All searchers start at the homebase — a selected node of the network. The network itself is not known in advance to the searchers. In fact, the searchers have no information about the network. (We note here that our main algorithmic result will be obtained in two stages: first we describe an algorithm that as an input receives the upper bound n of the size of the network and then we use it to obtain our main result, an algorithm that works without any a priori information about the network.) We assume that nodes are anonymous and searchers have identifiers. The edges incident to each node are marked with unique labels (port numbers) and because only partial grids are considered in this work (for a definition see Section 2.2) we assume that labels naturally reflect all possible directions for each edge (i.e., left, right, up and down).

For the searchers, we assume that they communicate *locally* by exchanging information when present at the same node. Our algorithm is stated as there existed *global* communication but it can be easily turned into required one with local communication as follows: we can designate one searcher called the *leader* who will be performing the following actions at the beginning of each move of the search strategy to be executed. First, the leader visits all nodes of the subgraph searched to date and gathers complete information about its structure and positions of all other searchers, then the leader computes the next move and finally visits all searchers to pass the information about the next move. Then, the move is performed by the agents.[‡]

Our algorithm is described for the synchronous model in which time is divided into steps, each step having the same unit length duration allowing each searcher to perform its local computations and slide along an edge if the searcher decides to move. We note that this assumption can be lifted and the algorithm can be easily restated to be asynchronous. Indeed, having one agent that is the leader one can simulate synchronous behavior of the agents in such a way that the leader waits for the completion of the current move of another searcher and then informs the searcher that is supposed to perform the next move, dictated by the search strategy, to start the move.

As to the memory model, our algorithm requires that the memory size of the searchers is polynomial in the size of the network, and we do not attempt to optimize this parameter.

2.2 Partial grid notation

We assume that a partial grid graph is embedded into two-dimensional Cartesian coordinate system, with a horizontal x-axis and vertical y-axis. For convenience, the homebase is located in a $(0, 0)$ position. We define a partial grid $G = (V, E)$ as a set of nodes V and edges E, where following conditions hold:

1. for every $v = v\,(x, y) \in V$, x and y are integer coordinates of the node v,

[‡]Note that the actions of the leader clearly contain a lot of excess work in terms of the number of moves it performs; since the criteria as time or cost (number of sliding moves) are out to scope of this work, we will leave the reader with such a simple leader implementation.

2. for any two adjacent nodes $v\,(x,y)$ and $v\,(x',y')$ the distance between (x,y) and (x',y') equals one (in Euclidean metric).

In this work n denotes an upper bound of the number of nodes of a partial grid, such that $\sqrt{n}$ is an integer.

Let us notice here that nodes at distance one in the grid are not necessarily neighbors in the graph, thus a partial grid defined in such a way can take all possible shapes, i.e., it can be a tree or a mesh, and it can contain holes. We note that some simple graphs can be embedded in various ways and for different embeddings our algorithm may perform differently.

Informally speaking, our algorithm will conduct a search by expanding the clean part of the graph from one 'checkpoint' to another. These checkpoints (defined formally later) will be subsets of nodes and their potential placements on the partial grid are dictated by a concept of a *frontier*. Take any $x = i\sqrt{n}$ for some integer i, $y = j\sqrt{n}$ for some integer j and take $i', j' \in \{0,1\}, i' \neq j'$. Then, the line segment with endpoints (x,y) and $(x+\sqrt{n}i', y+\sqrt{n}j')$ is called a *frontier* and denoted by $F\left((x,y),(x+\sqrt{n}i', y+\sqrt{n}j')\right)$. Whenever the endpoints of a frontier are clear from the context or not important we will omit them. The frontier $F\left((0,0),(\sqrt{n},0)\right)$ that contains the origin is called the *homebase frontier* and the set of all frontiers is denoted by $\mathcal{F}$. We will also divide frontiers into *vertical* and *horizontal* ones, where coordinates of two extreme nodes do not differ on first and second coordinate, respectively.

The subgraph induced by all nodes that belong to a frontier F is denoted by $G[F]$.

For $i \in \{1,\ldots,\sqrt{n}\}$ and some frontier $F = F\left((x,y),(x',y')\right)$, where $x \leq x'$ and $y \leq y'$, we define *an i-th rectangle of F*, denoted by $\mathcal{R}(F,i)$, as the rectangle with corner vertices $(x-i, y-i)$, $(x-i, y+i)$, $(x'+i, y'-i)$, $(x'+i, y'+i)$ if F is horizontal and as the rectangle with corner vertices $(x-i, y-i)$, $(x+i, y-i)$, $(x'-i, y'+i)$, $(x'+i, y'+i)$ if F is vertical.

Informally speaking, the two above concepts, namely frontiers and rectangles, provide a template on how the search may progress. However, due to the structure of a partial grid it may be possible that only certain nodes, but not all, that lie on a frontier are reached at some point of a search strategy. For this reason, our notation needs to be extended to subsets of nodes that lie on frontiers and the corresponding rectangles. Let $F = F\left((x,y),(x',y')\right)$ be some frontier. Any subset C of nodes of G that belong to F is called a *checkpoint*. *The 0-th expansion of a checkpoint C* is C itself and is denoted by $C\langle 0\rangle$. For $i \in \{1,\ldots,\sqrt{n}\}$ we define *an i-th expansion of C*, denoted by $C\langle i\rangle$, recursively as follows: the set $C\langle i\rangle$ consists of all nodes $v \notin C\langle 0\rangle \cup C\langle 1\rangle \cup \cdots \cup C\langle i-1\rangle$ for which there exists a node $u \in C\langle i-1\rangle$, such that there exists a path between v and u in the subgraph of G induced by nodes that lie on the rectangles $\mathcal{R}(F,0), \mathcal{R}(F,1), \ldots, \mathcal{R}(F,i)$. Define

$$C^{+}\langle i\rangle = C\langle 0\rangle \cup \ldots \cup C\langle i\rangle, \quad i \in \{0,\ldots,\sqrt{n}\}.$$

Informally, $C\langle i\rangle$ consists of only those nodes that belong to the rectangle $\mathcal{R}(F,i)$ that are connected to nodes of C by paths that lie 'inside' of $\mathcal{R}(F,i)$ — this definition captures the behavior of searchers (in our algorithm) that guard the nodes of C and 'expand' from C in all directions: then possible nodes that belong to any of the rectangles $\mathcal{R}(F,0), \mathcal{R}(F,1), \ldots, \mathcal{R}(F,i)$ but do not belong to $C^{+}\langle i\rangle$ will not be reached by the searchers. See Figure 1 for an exemplary checkpoint with its expansions.

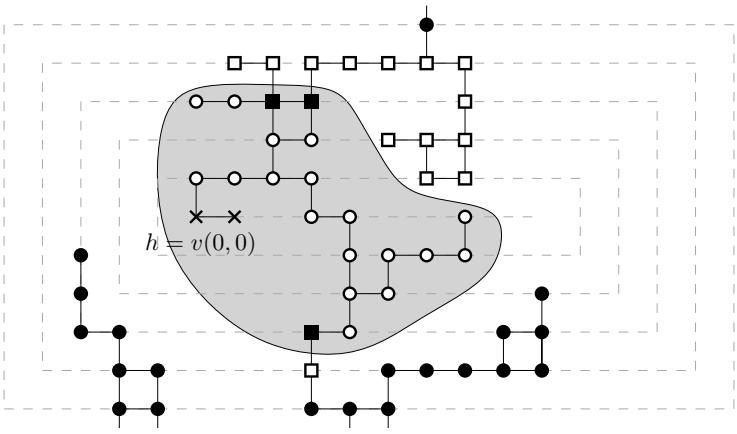

Figure 1: Exemplary expansions of a checkpoint C (here $\sqrt{n} = 9$); *crosses* denote $C = C\langle 0\rangle$, *gray area* covers nodes that belong to $C^+\langle 3\rangle$, *empty squares* denote nodes in $C\langle 4\rangle$ and *dark squares* denote the one that need to be guarded provided that the gray area consists of the clean nodes.

3 Lower bound

First note that a regular $\sqrt{n} \times \sqrt{n}$ grid requires $\Omega(\sqrt{n})$ searchers even in the offline setting [14], that is, when the network is know in advance and the searchers may decide on the location of the homebase. Therefore, our distributed algorithm is asymptotically optimal with respect to this worst case measure.

What we would also like to obtain is a lower bound expressed as a competitive ratio, which is defined as a maximized over all networks and all starting nodes ratio between number of searchers that a given algorithm uses and the search number that is optimal for a given network in an offline settings. In other words, for any distributed algorithm A, let $A(G,h)$ be the number of searchers that it uses to clean a network G starting from the homebase h and let $A(G) = \max_h A(G,h)$. We aim at proving that for *each* distributed algorithm A there exists an n-node partial grid network G such that $A(G)/\texttt{mcs}(G) = \Omega(\sqrt{n}/\log n)$.[§]

Define a class of partial grids

$$\mathcal{L} = \bigcup_{l \geq 0} \mathcal{L}_l,$$

where $\mathcal{L}_l$ for $l \geq 0$ is defined recursively as follows. We take $\mathcal{L}_0$ to contain one network that is a single node located at $(0,0)$. Then, in order to describe how $\mathcal{L}_{l+1}$ is obtained from $\mathcal{L}_l$, $l \geq 0$, we introduce an operation of *extending* $G \in \mathcal{L}_l$ *at* i, for $i \in \{0,\ldots,l\}$. In this operation, first take G and add $l+2$ new nodes located at coordinates:

$$(0, l+1), (1, l), \ldots, (j, l+1-j), \ldots, (l+1, 0).$$

[§]We remark that we define this competitive ratio by taking the worst case homebase for A and in the definition of $\texttt{mcs}(G)$ the most favorable homebase is selected. However, we note that this does not weaken the result of this section as, informally speaking, one may take two copies of each grid obtained in this section, rotate one copy by 180 degrees and merge the two copies at their homebases. Then, we obtain that for each choice of the homebase any algorithm is forced to use $\Omega(\sqrt{n})$ searchers for some grids since in one copy the search is conducted as in our following analysis.

Call these coordinates the $(l+1)$-*th diagonal*. For each $j \in \{0, \ldots, i\}$ add an edge connecting the nodes $v(j, l-j)$ and $v(j, l-j+1)$, and for each $j \in \{i, \ldots, l\}$ add an edge connecting the nodes $v(j, l-j)$ and $v(j+1, l-j)$. Then, obtain $\mathcal{L}_{l+1}$ as follows: initially take $\mathcal{L}_{l+1}$ to be empty and then for each $G \in \mathcal{L}_l$ and for each $i \in \{0, \ldots, l\}$, obtain a network G' by extending G at i and add G' to $\mathcal{L}_{l+1}$. Notice here that a graph constructed this way is not only a partial grid, but also a tree.

Figure 2 shows a network that was obtained from the corresponding network in $\mathcal{L}_7$ by extending it at 6.

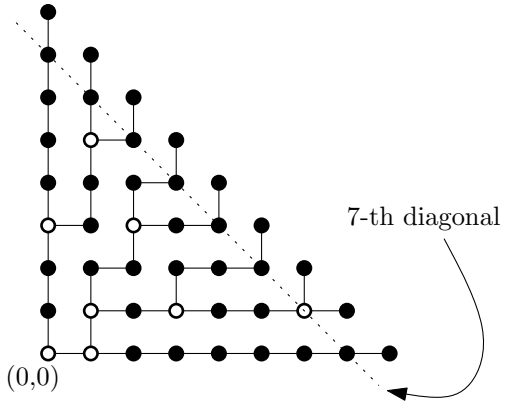

Figure 2: A network from $\mathcal{L}_8$

For a network $G \in \mathcal{L}_l$, $l \geq 0$, we define a *characteristic sequence of* G, $\sigma(G)$, as follows. If $l = 0$, then the characteristic sequence of G is empty. If $l > 0$, then take the network G' such that G has been obtained by extending G' at i. Then, $\sigma(G) = (\sigma(G'), v(i, l-i-1))$ is the characteristic sequence of G. Note that the characteristic sequence uniquely defines the corresponding network. The network introduced in Figure 2 has characteristic sequence $(v(0,0), v(1,0), v(1,1), v(0,3), v(3,1), v(2,3), v(1,5), v(6,1))$.

Lemma 3.1. *For any integer l and for each distributed algorithm A computing a connected monotone search strategy there exists $G \in \mathcal{L}_l$ such that for homebase $v(0,0)$ we have $A(G, v(0,0)) \geq (l+1)/2$.*

Proof. Consider any algorithm A producing a connected monotone search strategy. Run A for each network in $\mathcal{L}_l$ with the homebase $v(0,0)$. Note that for each network in $\mathcal{L}_l$, there exist distinct moves $m_1, \ldots, m_l$ such that till the beginning of move m_j, $j \in \{1, \ldots, l\}$, no node on the j-th diagonal has been occupied by a searcher and at the end of m_j some node $v(x_j, y_j)$ of the j-th diagonal is occupied by a searcher. Consider $G \in \mathcal{L}_l$ such that $\sigma(G) = (v(0,0), v(x_1, y_1), \ldots, v(x_{l-1}, y_{l-1}))$. Informally speaking, whenever the algorithm reaches for the first time a node $v(i, j-i)$ in the j-th diagonal, an *adversary* decides to extend at i the network explored so far, thus always forcing the situation that the first node reached on a diagonal is of degree three.

Note that at the beginning of move m_j, $j \in \{1, \ldots, l\}$, no node of the j-th diagonal has been reached by a searcher and the first j nodes of the characteristic sequence have been reached by searchers. Recall that G is a binary tree.

We analyze the explored part of $\mathcal{L}_l$ at the beginning of the move m_l. All edges incident to the leaves in $\mathcal{L}_l$ are contaminated at this point. On the other hand, all nodes of the

characteristic sequence have been visited by searchers till the end of the move $m_l - 1$. Therefore, the contaminated subgraph of $\mathcal{L}_l$ at this point is a collection of paths leading from nodes that are guarded to the leaves. Since there are $l + 1$ leaves in $\mathcal{L}_l$, there are $l + 1$ such paths, each such path needs to have a searcher placed at one of its endpoints (the one that is not a leaf in $\mathcal{L}_l$) and, by construction of $\mathcal{L}_l$, any searcher can be present on at most two such endpoints. Thus, at least $(l + 1)/2$ nodes need to be occupied by searchers, as required by the lemma. □

Theorem 3.1. *For each distributed algorithm A computing a connected monotone search strategy there exists an n-node network G with homebase h such that*

$$\frac{A(G,h)}{\texttt{mcs}(G)} = \Omega(\sqrt{n}/\log n).$$

Proof. Observe that each network G in $\mathcal{L}$ is a tree and therefore $\texttt{mcs}(G) = O(\log(n))$, $n = |V(G)|$ [1, 26]. The theorem follows hence from Lemma 3.1 and the fact that the length of characteristic sequence of each network in $\mathcal{L}_l$ is $\Omega(\sqrt{n})$. □

4 The algorithm

In this section we describe our algorithm that takes an upper bound on the size of the network as an input. Section 4.1 deals with the initialization performed at the beginning of the algorithm. Then, Section 4.2 introduces two procedures used by the algorithm and finally Section 4.3 states the main algorithm. After each move performed by searchers, each searcher that occupies a node that does not need to be guarded is said to be *free*. Each node that needs to be guarded is occupied by at least one searcher; if more searchers occupy such node then all of them except for one are also *free*. If, at some point, no node of the last expansion of some checkpoint need to be guarded, then we say that the expansion is *empty*.

4.1 Initialization

We start presenting our algorithm by describing initial conditions. Recall that the origin $v\,(0,0)$ of the two-dimensional xy coordinate system is situated in the homebase. The initial checkpoint C_0 is the set of nodes of the connected component of $G[F]$ that contains h, where F is the homebase frontier. Thus, initially $|C_0|$ searchers place themselves on all nodes of C_0 (note that the nodes of C_0 induce a path in G). See Figure 3 for an example.

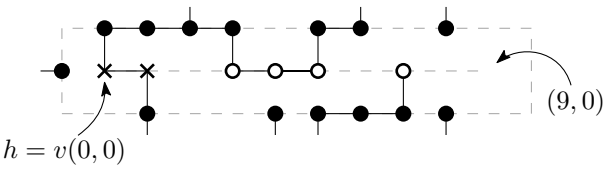

Figure 3: Exemplary initialization for $\sqrt{n} = 9$; *crosses* denote nodes belonging to the initial checkpoint C_0 and *empty circles* denote nodes that belong to the homebase frontier, but do not fall into C_0.

4.2 Procedures

4.2.1 Procedure `CleanExpansion`

We start with an informal description of the procedure. When a new checkpoint C has been reached, our search strategy 'expands' from C by successively cleaning subgraphs $G[C^+\langle i\rangle]$ for $i \in \{1, \ldots, \sqrt{n}\}$. Once all nodes in $C^+\langle i-1\rangle$ are clean for some $0 < i \leq \sqrt{n}$, the transition to reaching the state in which all nodes in $C^+\langle i\rangle$ are clean requires cleaning all nodes of the i-th expansion of C. This is done by calling for every guarded node u from $C^+\langle i-1\rangle$ a special procedure (`ModConnectedSearching`, described below), which cleans nodes which belong to $C\langle i\rangle$ and 'has access' to them from u. Procedure `CleanExpansion` performs this job by using $O(\sqrt{n})$ searchers.

For cleaning all nodes of the i-th expansion of C, provided that $G[C^+\langle i-1\rangle]$ is clean we will use a procedure from [5] which is more general but for our purposes can be stated as follows.

Theorem 4.1 ([5]). *Let F be any frontier and let G' be any connected grid with nodes lying on rectangles $\mathcal{R}(F,0), \mathcal{R}(F,1), \ldots, \mathcal{R}(F,i)$, $i \geq 0$. There exists a distributed procedure* `ConnectedSearching` *that, starting at an arbitrarily chosen homebase in G', clears G' in a connected and monotone way using $6i+4$ searchers.* □

Note that while using procedure `ConnectedSearching`, we will be cleaning a subgraph of $G[C\langle i\rangle]$ that is embedded into the entire partial grid and thus some nodes v of $G[C\langle i\rangle]$ have edges leading to neighbors that lie outside of $G[C\langle i\rangle]$. If such an edge is already clean, then no recontamination happens for the node v and moreover no searcher used by `ConnectedSearching` for the subgraph of $G[C\langle i\rangle]$ needs to stay at v. On the other hand, if such an edge is contaminated (and thus not reached yet by our search strategy), then v needs to be guarded and for that end we place an extra searcher on it that guards v during the remaining execution of `ConnectedSearching`. Note that in the latter case, the node v belongs to $\mathcal{R}(F,i)$, where F is the frontier that contains the nodes of C and therefore there exist $O(\sqrt{n})$ such nodes v. In other words, `ConnectedSearching` is called to clean a certain subgraph contained within $\mathcal{R}(F,i)$ and whenever a node on the rectangle $\mathcal{R}(F,i)$ has a contaminated edge leading outside of the rectangle $\mathcal{R}(F,i)$, then an extra searcher, no accommodated by `ConnectedSearching` in Theorem 4.1, is introduced to be left behind to guard v. The modification of `ConnectedSearching` that leaves behind a searcher on each such newly reached node of $\mathcal{R}(F,i)$ will be denoted by `ModConnectedSearching`. Note that this procedure is invoked for every guarded node from $C^+\langle i-1\rangle$ in order to clean $C\langle i\rangle$, see Figure 4 for an example. It follows that it is enough to provide as an input to `ModConnectedSearching`: a node in $C^+\langle i-1\rangle$ that plays the role of homebase for `ModConnectedSearching`, the frontier F and i. We note that each checkpoint constructed in our final algorithm is obtained as follows: some frontier F is selected and then a checkpoint C is created as some set of nodes that belong to F; thus we assume that with C is associated such a unique frontier F.

Thus, this approach guarantees us using at most $6i+4$ searchers to clean $G[C\langle i\rangle]$ and $2\sqrt{n}+8i$ searchers for guarding nodes laying on $\mathcal{R}(F,i)$, which will be analyzed in more details in Section 5.

To summary, we give a formal statement of our procedure.

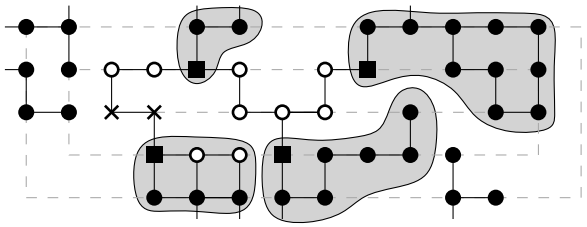

Figure 4: Example of an execution of procedure `CleanExpansion`; *crosses* denote $C = C\langle 0\rangle$, *empty circles* denote nodes that belong to $C^+\langle 1\rangle$, *dark squares* denote the one that belongs to $C^+\langle 1\rangle$ and for which procedure `ModConnectedSearching` is invoked, *gray areas* show nodes that will be cleaned in four calls of `ModConnectedSearching` in order to clean $C\langle 2\rangle$. Note that there are some nodes that belongs to $C^+\langle 1\rangle$ and are guarded at first, but after one of the calls of `ModConnectedSearching` there is no need to guard them any more, so the procedure is not invoked for them.

Procedure `CleanExpansion`

Input: An expansion $C\langle i-1\rangle$ with C contained in the frontier F, $i \geq 1$.
Result: Cleaning all nodes of $C\langle i\rangle$.
while there exists a node $v \in C\langle i-1\rangle$ with a contaminated neighbor u in $C\langle i\rangle$ **do**
Place $6i+4$ free searchers on v.
Call `ModConnectedSearching` for v as the homebase, frontier F and integer i.

4.2.2 Procedure `UpgradeCheckpoints`

Our algorithm maintains a collection $\mathcal{C}$ of currently used checkpoints. Note that $\mathcal{R}(F, \sqrt{n})$, where F is some frontier, contains 10 frontiers. Thus, reaching the $\sqrt{n}$-th expansion of a checkpoint of F provides a possibility of creating one new checkpoint for each of the above frontiers. Procedure `UpgradeCheckpoints` generates these new checkpoints, adds them to $\mathcal{C}$ and removes C from $\mathcal{C}$. Also, if it happens that some newly constructed checkpoint belongs to the same frontier as some existing checkpoint in $\mathcal{C}$ and no expansion for the existing one has been performed yet, then both checkpoints are merged into one. Finally, any checkpoint in $\mathcal{C}$, whose lastly performed expansion is empty, is removed from $\mathcal{C}$. We remark that only procedure `UpgradeCheckpoints` modifies the collection of checkpoints $\mathcal{C}$ and this procedure performs no cleaning moves.

Thus, to summary, the 'lifetime' of a checkpoint is as follows. A newly created checkpoint C may change by getting more nodes only before its 1-st expansion occurs. This happens if for another checkpoint its $\sqrt{n}$-th expansion is performed and as a result some new nodes that belong to the same frontier as C become guarded. Once the 1-st expansion of C is performed, the checkpoint will remain in the collection $\mathcal{C}$ and possibly more expansions of C are made (in total at most $\sqrt{n}$ expansion are possible for each checkpoint). Finally, C may disappear from $\mathcal{C}$ in two ways: either some expansion of C becomes empty (then C is not removed from $\mathcal{C}$ right away but during the subsequent call to `UpgradeCheckpoints`), or C reaches its $\sqrt{n}$-th expansion and procedure `UpgradeCheckpoints` is called for C (in which case C possibly 'gives birth' to new checkpoints during the execution of `UpgradeCheckpoints`).

Procedure `UpgradeCheckpoints`

Input: $C\langle\sqrt{n}\rangle$ and the collection of all checkpoints $\mathcal{C}$
Result: Updated collection $\mathcal{C}$
  $\mathcal{C} \leftarrow \mathcal{C} \setminus \{C\}$
  $\mathcal{C}_{new} \leftarrow \emptyset$
  for each frontier F on $\sqrt{n}$-th rectangle of the frontier containing C **do**
    Let C' consist of all guarded nodes in F.
    If $C' \neq \emptyset$, then $\mathcal{C}_{new} \leftarrow \mathcal{C}_{new} \cup \{C'\}$.
  for each C'' in $\mathcal{C}$ **do**
    if there exists $C' \in \mathcal{C}_{new}$ that is a subset of the same frontier as C'' **then**
      if C'' is in 0-th expansion **then**
        $\mathcal{C}_{new} \leftarrow \mathcal{C}_{new} \setminus \{C'\}$
        Replace C'' with $C'' \cup C'$ in $\mathcal{C}$.
  $\mathcal{C} \leftarrow \mathcal{C} \cup \mathcal{C}_{new}$
  for each C in $\mathcal{C}$ **do**
    if no node in the last expansion of C is guarded **then**
      $\mathcal{C} \leftarrow \mathcal{C} \setminus \{C\}$

4.3 Procedure `GridSearching`

`GridSearching` is the main algorithm, whose aim is to clear the entire partial grid G in the connected and monotone way. We start with an informal introduction of the algorithm. The search strategy it produces is divided into phases, which formally will be defined in the next chapter. In each step of the algorithm, the checkpoint with the highest number of nodes that need to be guarded is being chosen and the next expansion is being made on it. When one of the checkpoints reaches its $\sqrt{n}$-th expansion, then the current phase ends and the procedure `UpgradeCheckpoints` is being invoked. Thus, the division of the search strategy into phases is dictated by consecutive calls to procedure `UpgradeCheckpoints`. For an expansion C, in the pseudocode below we write $\delta(C)$ to refer to the set of nodes that belong to the last expansion of C and need to be guarded at a given point.

Procedure `GridSearching`

Input: An integer n providing an upper bound on the size of the partial grid G.
Result: A monotone connected search strategy for G.
  Perform the initialization (see Section 4.1).
  while G is not clean **do**
    while no checkpoint has reached its $\sqrt{n}$-th expansion **do**
      Let $C_{max} \in \mathcal{C}$ be such that $\delta(C_{max}) \geq \delta(C)$ for each $C \in \mathcal{C}$.
      Let i be the number of expansions of C_{max} performed so far.
      Invoke `CleanExpansion` for $C_{max}\langle i\rangle$.
    Invoke `UpgradeCheckpoints` for $C_{max}\langle\sqrt{n}\rangle$ and $\mathcal{C}$.

We close this chapter with giving examples of first three expansions of some checkpoint C, see Figure 5, and showing how our algorithm clears an exemplary partial grid network, see Figure 6.

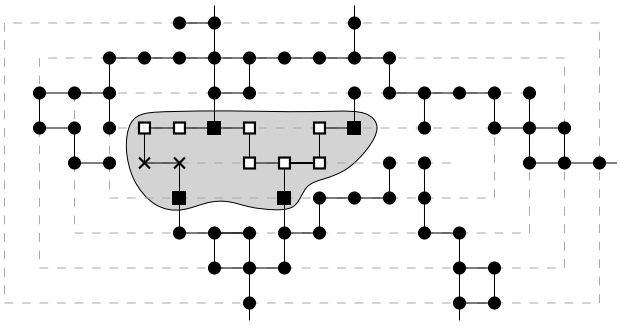

(a) First expansion of C.

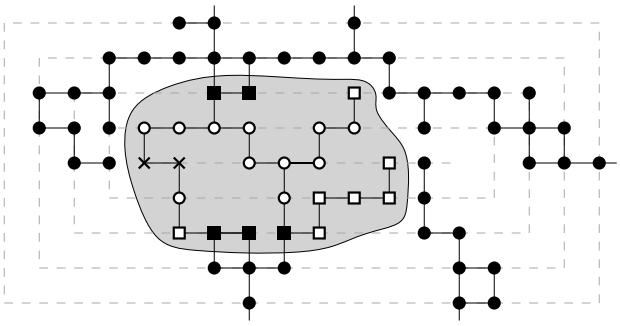

(b) Second expansion of C.

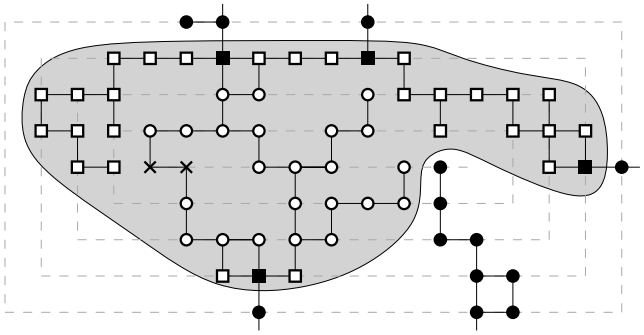

(c) Third expansion of C.

Figure 5: First three expansions for some checkpoint C (here $\sqrt{n} = 9$); *crosses* denote $C = C\langle 0 \rangle$, *empty circles* denote nodes cleaned in previous expansions; *squares* denote nodes explored in current expansion; *dark circles* are nodes not reached yet by the searchers; and *dark squares* denote nodes that need to be guarded at the end of current expansion. *Gray areas* show the clean part of the graph, i.e., $C^+\langle i \rangle$ for $i \in \{1, 2, 3\}$.

5 Analysis of the algorithm

By a *step of the algorithm*, or simply a *step*, we mean all searching moves performed during a single iteration of the internal 'while' loop of procedure `GridSearching`. Thus, one step of the algorithm includes all moves produced by one call to procedure `CleanExpansion`. A *phase* of an algorithm consists of all its steps between two consecutive calls to procedure `UpgradeCheckpoints`. Note that phases may differ with respect to the number of steps they are made of.

We say that a checkpoint is *present* in a given phase if its last expansion is not empty at the beginning of this phase, i.e., if this checkpoint belongs to $\mathcal{C}$ at the beginning of the phase. Similarly, a checkpoint is *present* in a given step if it is present in the phase to which the step belongs to. Thus, in particular, a checkpoint is present in none or in all steps of a given phase. Note that some checkpoints may have empty expansions during a part of a the phase, but they still remain present to the end of the phase; this assumption is made to simplify the analysis of the algorithm.

For the purposes of the next definition we say that, for a checkpoint C, a node v and a step t, the checkpoint C *owns* v *in step* t if:

- v needs to be guarded at the beginning of step t and v belongs to the last expansion of C performed till the end of step $t - 1$, and
- either this last expansion of C occurred in step $t - 1$, or C owns v in step $t - 1$.

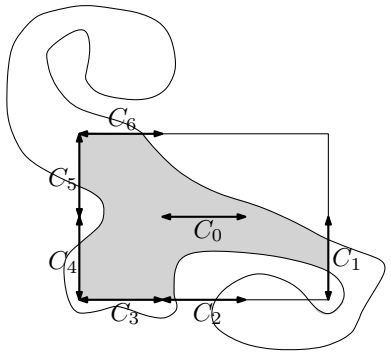

(a) At the end of the first phase C_0 (initial checkpoint) reaches its $\sqrt{n}$-th expansion. Procedure `UpgradeCheckpoints` creates 6 new checkpoints and removes C_0 from $\mathcal{C}$, i.e. $\mathcal{C} = \{C_1, C_2, C_3, C_4, C_5, C_6\}$.

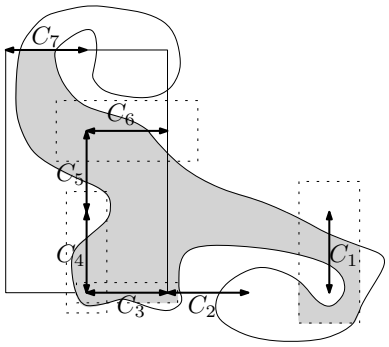

(b) $\sqrt{n}$-th expansion of C_5 ends the second phase. Checkpoints C_4 and C_6 are removed from $\mathcal{C}$ because (in our example) there is no need to guard any node on theirs expansions; $\mathcal{C} = \{C_1, C_2, C_3, C_7\}$.

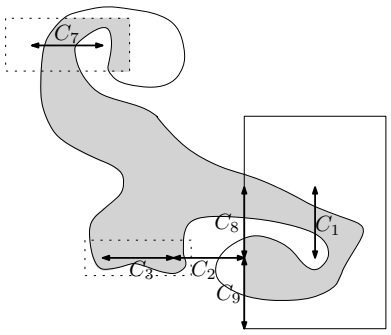

(c) C_1 ends the third phase. Notice that a new checkpoint C_8 emerged 'inside' already cleaned area by C_0; C_2 is removed from $\mathcal{C}$ even if $\sqrt{n}$-th expansion has not been reached but its last expansion has no nodes to be guarded; $\mathcal{C} = \{C_7, C_8, C_9\}$.

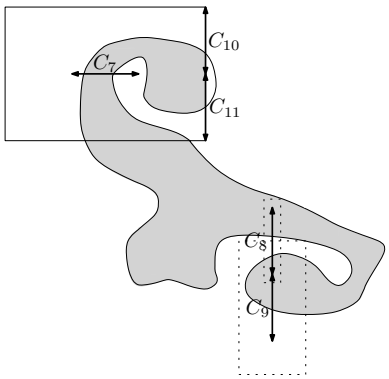

(d) C_7 ends the fourth phase. Notice that a new checkpoint C_{11} emerged on an edge of C_5's $\sqrt{n}$-th expansion, but it could not be created in the second phase because then there was no access to the contaminated part; $\mathcal{C} = \{C_{10}, C_{11}\}$.

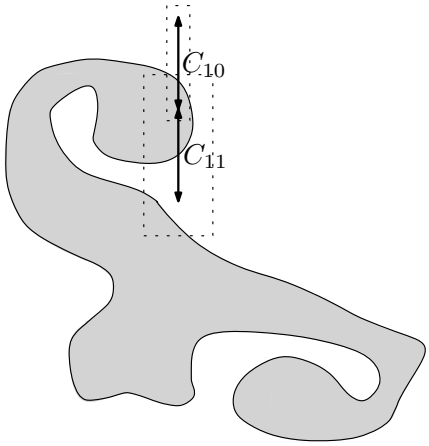

(e) Last phase, in which the rest of the graph is cleaned.

Figure 6: Clearing an exemplary partial grid by procedure `GridSearching`; *gray areas* denote the clean part, *arrows* denote frontiers on which the marked checkpoints lie, *dotted rectangles* around checkpoints denote their current expansions and *solid rectangles* denote the $\sqrt{n}$-th expansions, which end phases.

(Intuitively, if a node v is reached by searchers in a step in which an expansion of C occurred, then C owns v as long as v is guarded.) Given a checkpoint C present in a step t, we write $\mathcal{E}(C,t)$ to denote the set of nodes that C owns in step t. The *weight of a checkpoint* C present in a step t is $\omega_t(C) = |\mathcal{E}(C,t)|$. Note that each guarded node is owned by exactly one checkpoint and hence, for a step t, the sum of weights of all checkpoints present in step t equals the number of nodes that need to be guarded.

If a checkpoint C is not present in a step t, then we take $\omega_t(C) = 0$. The checkpoint C_{max} selected in a step t (see the pseudocode of Procedure `GridSearching`) is called *active in step* t, or simply *active* if the step is clear from the context or not important. All other checkpoints present in this step are called *inactive*. We define an *active interval* of a checkpoint C to be a maximal interval $[t', t'']$ such that C is active in all steps $t \in \{t', \ldots, t''\}$.

5.1 Single phase analysis — how weights of checkpoints evolve

We now prove lemmas that characterize how the weight of a checkpoint changes over time — see Figure 7 for an exemplary lifeline of a checkpoint. Informally, the weight of a checkpoint C does not grow in intervals in which C is inactive (Lemma 5.1). Also, the weight of C at the end of an active interval is not greater than at the beginning of it (Lemma 5.2); however, no upper bounds except for the trivial one of $O(\sqrt{n})$ can be concluded for the weight of C inside its active interval.

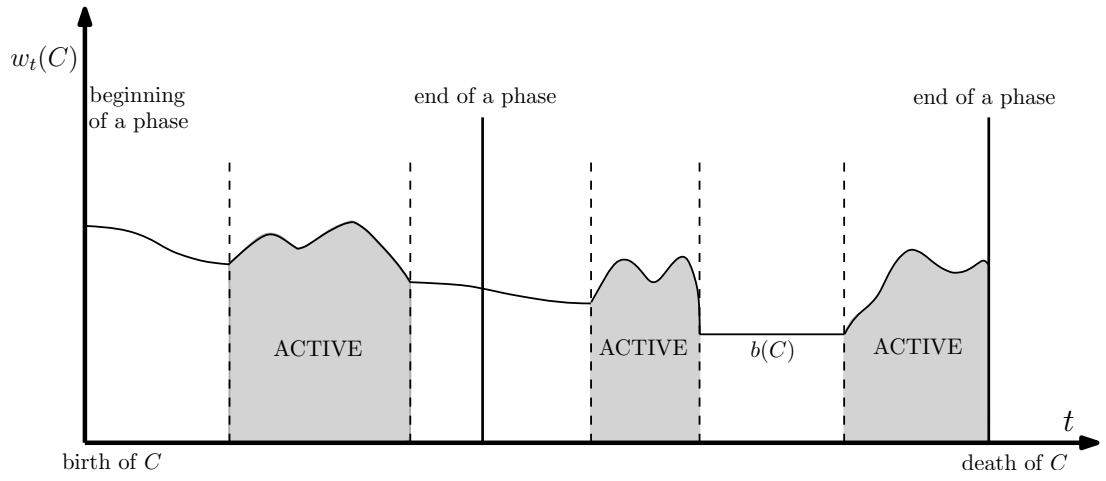

Figure 7: Exemplary lifeline of a checkpoint C.

Lemma 5.1. *If a checkpoint C is present and inactive in a step t, then $\omega_{t+1}(C) \leq \omega_t(C)$.*

Proof. It follows directly from the definitions and procedure `CleanExpansion` that the only checkpoint on which an expansion is performed during execution of `CleanExpansion` is the active one. The weight of an inactive checkpoint C can change only in the situation, where the active checkpoint in a step t expands on some nodes that belong to the last expansion of C. In other words, the weight of C may decrease if C contains in step t nodes that are added to the active checkpoint in step $t+1$. Thus, if t is not the last step of a phase, then the proof is completed.

If t is the last step of some phase, then apart from procedure `CleanExpansion`, procedure `UpgradeCheckpoints` is being invoked, which affects C in two situations:

- there exists a step t' in the phase that ends such that $\omega_{t'}(C) = 0$. Then, because C can not be expanded during steps $t', \ldots, t$ of the phase, we get directly that $\omega_{t+1}(C) = \omega_t(C) = 0$.
- C is in its 0-th expansion and a new checkpoint is placed on the same frontier, which implies that C is not present in step $t+1$ and thus $\omega_{t+1}(C) = 0$.

Thus, in all cases we obtain that $\omega_{t+1}(C) \leq \omega_t(C)$. □

We next observe that, informally speaking, once a checkpoint becomes active, it remains active until either the phase ends or its weight decreases. Note that a checkpoint that is active in the last step of the phase is not present in the first step of the next phase, i.e. its weight is then zero, which allows us to state the lemma as follows:

Lemma 5.2. *Let C be a checkpoint and let $[t', t'']$ be an active interval of C. Then, $\omega_{t''+1}(C) \leq \omega_{t'}(C)$.*

Proof. Obviously, t' and t'' must belong to the same phase, because at the end of each phase the active checkpoint is removed from $\mathcal{C}$, i.e., it is no longer present in the next phase.

If t'' is the last step of the phase then the lemma follows, because $\omega_{t''+1}(C) = 0 \leq \omega_{t'}(C)$.

We will now prove that lemma holds when t'' is not the last step of the phase. Let us suppose for a contradiction that $\omega_{t''+1}(C) > \omega_{t'}(C)$. From the assumptions of the lemma and definition of an active interval we get that C is not the active checkpoint in step $t''+1$. Because we are still in the same phase, it means that there must exist a checkpoint C^* such that $\omega_{t''+1}(C^*) \geq \omega_{t''+1}(C)$. Moreover from Lemma 5.1 we know, that because C^* was inactive from step t' to t'', it holds $\omega_{t'}(C^*) \geq \omega_{t''}(C^*) \geq \omega_{t''+1}(C^*)$. This gives us

$$\omega_{t'}(C^*) \geq \omega_{t''+1}(C^*) \geq \omega_{t''+1}(C) > \omega_{t'}(C),$$

which is in a contradiction to the assumption, that C is the active checkpoint in step t'. □

Remark 5.1. *Let C be a checkpoint and let $[t', t'']$ be an active interval of C. For every step $t \in \{t', \ldots, t''\}$ it holds $\omega_t(C) \geq \omega_{t''+1}(C)$.* □

We now conclude from the two previous lemmas that the weight of each checkpoint is not greater at the end of a phase than at the beginning of the phase.

Lemma 5.3. *Suppose that a phase starts in step t' and ends in step t''. For each checkpoint C present in this phase it holds $\omega_{t''+1}(C) \leq \omega_{t'}(C)$.*

Proof. Each checkpoint C can be active or inactive in different steps during the whole phase. If in some step $t \in \{t', \ldots, t''\}$ a checkpoint C is inactive then from Lemma 5.1 we have that the weight of it will not increase, i.e., $\omega_t(C) \geq \omega_{t+1}(C)$. On the other hand, Lemma 5.2 guarantees us, that the weight of an active checkpoint can not be greater at the end of any its active interval than at the beginning. Since end of a phase is the end of some active interval, this finishes the proof. □

Remark 5.2. *Suppose that a phase ends in a step t and the next one ends in a step t'. If a checkpoint C is inactive (but present) in steps t and t', then $\omega_{t'}(C) \leq \omega_t(C)$.* □

5.2 How many nodes are explored by a checkpoint?

Define a *bottleneck* of a checkpoint C, denoted by $b(C)$ to be its minimum weight taken over all steps in which C was present. (Note that a checkpoint may be present in many consecutive phases.)

Suppose that a node v has been reached by a searcher for the first time in a step t. Let C be the active checkpoint in step t. We say that v has been *explored by* C.

If an expansion of an active checkpoint C reaches in a step t a node u already explored by some checkpoint C', then in most situations u does not need to be guarded. However there might occur a "corner situation", when u still needs to be guarded in order to avoid contamination. In such case, the algorithm clearly needs one searcher on u to guard it and so it is counted in our analysis due to the 'ownership' relation used in the definition of the weight of a checkpoint.

The next lemma states a lower bound on the number of nodes explored by a checkpoint reaching its last expansion.

Lemma 5.4. *Suppose that a phase ends in a step t. Let C be the active checkpoint in step t. The number of nodes explored by C in all steps is at least $b(C)\sqrt{n}$.*

Proof. First let us make a remark that nodes can be only explored by C during execution of procedure `CleanExpansion` that took C as an input, i.e., when C is active. Let us denote by S the set of all nodes explored by C.

Because C is active in the last step of the phase, it had to be active in exactly $\sqrt{n}$ steps in total, which can be concluded in several past phases. Let $t_1, t_2, \ldots, t_{\sqrt{n}} = t$ be all steps in which C is active. Note that

$$\bigcup_{i=1}^{\sqrt{n}} \mathcal{E}(C, t_i) \subseteq S$$

and $\mathcal{E}(C, t_i) \cap \mathcal{E}(C, t_j) = \emptyset$ for $i \neq j$. The latter follows directly from the fact that nodes in $\mathcal{E}(C, t_i)$ and $\mathcal{E}(C, t_j)$ belong to different rectangles of the frontier containing C for $i \neq j$. (Recall that $|\mathcal{E}(C, t)| = \omega_t(C)$ for each step t.) Also from definition of the bottleneck, we get that $b(C) \leq \omega_{t_i}(C)$ for each $i \in \{1, \ldots, \sqrt{n}\}$ and hence we conclude that:

$$|S| \geq \sum_{i=1}^{\sqrt{n}} \omega_{t_i}(C) \geq b(C)\sqrt{n}.$$

□

We now give an upper bound on the weight of each inactive checkpoint at the end of a phase.

Lemma 5.5. *Suppose that a phase ends in a step t. Let $C_1, \ldots, C_l$ be all checkpoints present in this phase, where C_1 is the active checkpoint in step t. Then, $b(C_1) \geq \omega_t(C_j)$ for each $j \in \{2, \ldots, l\}$.*

Proof. Let us denote by t' the last step in which $\omega_{t'}(C_1) = b(C_1)$. If $t' = t$ then the lemma follows strictly from the definition of an active checkpoint. We will now prove that lemma stands also when $t' < t$.

Suppose that t' and t do not belong to the same active interval of C_1. From the Remark 5.1 we know that $\omega_{t''}(C_1) = b(C_1)$ occurs for some t'' that does not belong to an active interval. Moreover from Lemma 5.2 we get that every next active interval will need to start and finish on the same weight as the bottleneck, which is in contradiction that t' is the last step when $b(C_1)$ occurred.

Hence there must exist an active interval of C_1 that contains both t' and t. Then, we get from Lemma 5.1 and the fact that C_1 is active in step t':

$$\omega_t(C_j) \leq \omega_{t'}(C_j) \leq \omega_{t'}(C_1) = b(C_1), \quad j \in \{2, \ldots, l\},$$

which finishes our proof. □

Let us introduce a relation $\prec$ on a set of checkpoints. Whenever $C \prec C'$, we say that C is a *predecessor* of C' and C' is a *successor* of C. We stress out that the construction depends on the execution of the algorithm, namely only checkpoints that appear in some step are considered, and the division of the steps into phases shapes the relation. More precisely, the relation is defined only for checkpoints added to the set $\mathcal{C}$ during all executions of procedure `UpgradeCheckpoints`. To construct the relation we iterate over the consecutive phases of the algorithm. Initially the relation is empty and once the construction is done for each phase smaller than i, we perform the following for the phase i. Let C be the active checkpoint in the last step of phase i. Let $C_1, \ldots, C_l$ be all checkpoints, different than C, that have no successors so far and were added to $\mathcal{C}$ till the end of phase $i-1$. Then, let $C_j \prec C$ for each $j \in \{1, \ldots, l\}$.

An important property of our algorithm is that each checkpoint may have only constant number of predecessors:

Lemma 5.6. *Each checkpoint has at most* 10 *predecessors.*

Proof. A checkpoint C can only once be active in the last step of some phase i, because after that it will not be present in any later phases. At the end of phase i the only checkpoints that do not have any successors are the one that were constructed by the procedure `UpgradeCheckpoints` at the end of phase $i-1$. There are at most 10 such checkpoints. □

5.3 The algorithm uses $O(\sqrt{n})$ searchers in total

We now bound the total weight of all checkpoints at the end of each phase — note that this bounds the total number of searchers used for guarding at the end of a phase. A high level intuition behind the proof of Lemma 5.7 is as follows. Since, due to Lemma 5.4, each checkpoint C that is active in the last step of a phase explores at least $b(C)\sqrt{n}$ nodes in total. Therefore, the sum of bottlenecks of all such checkpoints C cannot exceed $\sqrt{n}$. Moreover, C can have at most 10 predecessors and hence the sum of weights of those predecessors is bounded by $10b(C)$ according to Lemma 5.5. Since each checkpoint (except the one that is active in the last step of a given phase) is a predecessor of some checkpoint that is active in the last step of some phase, we bound the sum of all weights of all such checkpoints present in a given phase by $10\sqrt{n}$.

Lemma 5.7. *Suppose that $C_1, \ldots, C_l$ are all checkpoints of a phase that ends in step t, where C_1 is active in step t. Then,*

$$\sum_{i=1}^{l} \omega_t(C_i) \leq \omega_t(C_1) + 10\sqrt{n}.$$

Proof. Suppose that a phase j ends in step t. Let t_i be the last step of phase i and let C_i^0 be the active checkpoint in step t_i for each $i \in \{0, \ldots, j\}$. We denote by s the number of nodes visited by searchers till the end of step t. From Lemma 5.4 and the fact that the number of all nodes n is at least s we have:

$$n \geq s \geq \sum_{i=0}^{j} b(C_i^0)\sqrt{n} \quad \Rightarrow \quad 10\sqrt{n} \geq 10 \sum_{i=0}^{j} b(C_i^0). \tag{1}$$

From Lemma 5.6 we have that the checkpoints $C_0^0, \ldots, C_j^0$ can have at most 10 predecessors. From the definition, they are constructed (i.e., added to collection $\mathcal{C}$ during execution of procedure `UpgradeCheckpoints`) at the beginning of the first step of a phase at the end of which their successor is active. Let us denote by $C_i^1, \ldots, C_i^{l_i}$, $0 \leq l_i \leq 10$, the predecessors of C_i^0 for each $i \in \{0, \ldots, j\}$ (by $l_i = 0$ we understand that C_i^0 has no predecessors). From Lemma 5.5 we have:

$$\sum_{k=1}^{l_i} \omega_{t_i}(C_i^k) \leq 10b(C_i^0), \quad i \in \{0, \ldots, j\}. \tag{2}$$

Remark 5.2 assures us that weights of inactive checkpoints will not be greater at the end of the next phase than they are in the last step of current phase:

$$\omega_t(C_i^k) = \omega_{t_j}(C_i^k) \leq \omega_{t_{j-1}}(C_i^k) \leq \cdots \leq \omega_{t_i}(C_i^k), \quad i \in \{0, \ldots, j\}; \quad k \in \{1, \ldots, l_i\}. \tag{3}$$

Because

$$\{C_1, \ldots, C_l\} \subseteq \{C_j^0\} \cup \left\{ C_i^k \;\middle|\; k \in \{1, \ldots, l_i\}, i \in \{0, \ldots, j\} \right\},$$

we can conclude from Equations (3), (2) and (1) (in this order) that:

$$\begin{aligned}
\sum_{i=1}^{l} \omega_t(C_i) &\leq \omega_t(C_j^0) + \sum_{i=0}^{j} \sum_{k=1}^{l_i} \omega_t(C_i^k) \\
&\leq \omega_t(C_j^0) + \sum_{i=0}^{j} \sum_{k=1}^{l_i} \omega_{t_i}(C_i^k) \\
&\leq \omega_t(C_j^0) + \sum_{i=0}^{j} 10b(C_i^0) \\
&\leq \omega_t(C_j^0) + 10\sqrt{n}.
\end{aligned}$$

□

Theorem 5.1. *Given an upper bound n of the size of the network as an input, the algorithm* `GridSearching` *clears in a connected and monotone way any unknown underlying partial grid network using $O(\sqrt{n})$ searchers.*

Proof. We will bound the number of searchers s used by a single call to procedure `CleanExpansion` and the total number of searchers s' used for guarding at the end of each step of the algorithm. Note that $s + s'$ bounds the total number of searchers used by `GridSearching`.

We first analyze procedure `CleanExpansion` to give an upper bound on s. We can divide searchers into three groups: *explorers*, *cleaners* and *guards*. Suppose that procedure `CleanExpansion` performs i-th expansion of a checkpoint C_{max}. Denote by F_{max} the frontier that contains the nodes in C_{max}. All searchers located at nodes on the $(i-1)$-th rectangle of F_{max} that need to be occupied in order to avoid recontamination at the beginning of the call to procedure `CleanExpansion` are named to be guards. The explorers and cleaners are used by algorithm `ModConnectedSearching` called during the execution of procedure `CleanExpansion`. Each time `ModConnectedSearching` reaches a node v on the i-th rectangle of F_{max} such that v needs to be guarded, the searcher used for guarding v is called an explorer. The searchers used in `ModConnectedSearching` that mimic the movements of searchers in algorithm `ConnectedSearching` are the cleaners. We point out that we do not alter here the behavior of `CleanExpansion` and `ModConnectedSearching` but just assign one of the three categories to each searcher they use. Informally speaking, when explorers protect nodes lying on the i-th rectangle and the guards protect the ones lying on the $(i-1)$-th rectangle of F_{max}, cleaners clean nodes inside the i-th rectangle of F_{max} (i.e., the remaining nodes of the i-th expansion of C_{max}).

The fact that each rectangle of a frontier contains at most $10\sqrt{n}$ nodes and Theorem 4.1 give that:

$$\begin{aligned} \text{number of explorers} &\leq 10\sqrt{n}, \\ \text{number of cleaners} &\leq 6\sqrt{n}+4. \end{aligned}$$

Thus,

$$s \leq 16\sqrt{n}+4.$$

The guards used to protect nodes lying on the $(i-1)$-th rectangle are accounted for during the estimation of s' below.

We now bound the maximal number of searchers used for guarding at the end of each step t of our search strategy, which we denote by g_t. It is easy to see that $g_t \leq 10\sqrt{n}$ if t belongs to phase 0.

Let us now take any step t that belongs to an i-th phase, where $i > 0$ and denote by t' the last step of the phase $i-1$ and by C the active checkpoint in step t'. From Lemma 5.7 we know that $g_{t'} \leq \omega_{t'}(C) + 10\sqrt{n} \leq 20\sqrt{n}$. The latter inequality follows from the fact that all nodes in $\mathcal{E}(C, t')$ belong to the j-th rectangle of the frontier that contains C, $j \leq \sqrt{n}$, and the number of nodes in this rectangle is at most $10\sqrt{n}$.

We know now that every phase starts with at most $20\sqrt{n}$ guards. If t is the first step of an active interval of some checkpoint, then by Lemmas 5.1 and 5.2 we have that $g_t \leq g_{t'} \leq 20\sqrt{n}$. But if t is a step inside some active interval, then an active checkpoint can reach at most $10\sqrt{n}$ new nodes that need to be guarded. Because in one step only one checkpoint can be active that leads us to conclusion that for every step t we have $g_t \leq 30\sqrt{n}$. Therefore, we obtain that $s' \leq 30\sqrt{n}$.

Thus, we obtain $s + s' \leq 46\sqrt{n} + 4 = O(\sqrt{n})$ as required. □

6 Unknown size of the graph

The algorithm we have described needs to know an upper bound of the size of the underlying partial grid network G. In this section we design a procedure called `ModGridSearching` that performs the search using $O(\sqrt{n})$ searchers and having no prior information on the network. The procedure is based on a standard technique: guessing an upper bound on n by doubling potential estimate each time.

The procedure `ModGridSearching` is composed of a certain number of m *rounds*. In round i, procedure `GridSearching` first introduces $c\sqrt{2^i}$ new searchers called *i-th team*, where c is a constant from the asymptotic notation in Theorem 5.1. Then, a call to `GridSearching` is made, where procedure `GridSearching` is using only the searchers of the i-th team. The outcome can be twofold. The procedure may succeed in searching the entire graph and in such case the i-th round is the last one and `ModGridSearching` is completed, or the procedure may encounter a situation in which it would be forced to use more than $c\sqrt{2^i}$ searchers to continue. In such case `GridSearching` stops, the i-th round ends and the $(i+1)$-th round will follow. Once the i-th round is completed, the searchers of the i-th team stay idle indefinitely. We point out that during the execution of an i-th round, $i > 1$, procedure `GridSearching` using the searchers of the i-th team is ignoring the fact that the network may be partially clear as a result of the work done in previous rounds. Moreover, the searchers of j-th team for each $j < i$ are not used and thus also ignored during i-th round.

We close this section by giving an upper bound of the number of searchers that need to be used in the presented modified version of our algorithm.

Theorem 6.1. *There exists a distributed algorithm that clears (starting at an arbitrary homebase) in a connected and monotone way any unknown underlying partial grid network using $O(\sqrt{n})$ searchers. The algorithm receives no prior information on the network.*

Proof. Let n be the number of nodes of the partial grid network, which is unknown to our procedure. The number of rounds m fulfills $2^{m-1} < n \leq 2^m$, i.e. $m = \lceil \log_2 n \rceil$. At the end of i-th round, $c\sqrt{2^i}$ searchers need to stay in their last positions till the end of our procedure and are not used in subsequent rounds. This means that the total number of searcher s is upper bounded by a sum of searchers used in every round:

$$\begin{aligned} s &\leq c\sqrt{2} + c\sqrt{2^2} + \cdots + c\sqrt{2^{\lceil \log_2 n \rceil}} = c \sum_{j=1}^{\lceil \log_2 n \rceil} \left(\sqrt{2}\right)^j \\ &= \sqrt{2}c\frac{1-\sqrt{2}^{\lceil \log_2 n \rceil}}{1-\sqrt{2}} = \frac{\sqrt{2}c}{\sqrt{2}-1}\left(\sqrt{2^{\lceil \log_2 n \rceil}} - 1\right). \end{aligned}$$

Because $\sqrt{n} \leq \sqrt{2^{\lceil \log_2 n \rceil}} < \sqrt{2n}$, we conclude

$$s < \frac{\sqrt{2}c}{\sqrt{2}-1}\left(\sqrt{2n} - 1\right) \quad \Rightarrow \quad s = O(\sqrt{n}).$$

□

7 Conclusions

7.1 Graph-theoretic modeling

In this section we consider a searching scenario for two-dimensional environment and its modeling via graph theory. This provides another motivation for the distributed graph searching model used in this work. Consider a continuous search problem in which k searchers initially placed at the same location need to capture the fugitive hiding in an arbitrary polygon that possibly has holes. The polygon is not known a priori to the searchers. The fugitive is considered captured in time t when it is located at distance at most r from some searcher at time point t. (The r can be related to physical dimensions of searchers and/or their visibility range, etc.)

Consider the following transition from the above continuous searching problem of a polygon to a discrete one. Overlap the coordinate system with the polygon in such a way that the origin coincides with the original placement of the searchers. Then, place nodes on all points with coordinates being multiples of r and lying in the polygon. Connect two nodes with an edge if the edge is contained in the polygon. In this way we obtain a partial grid network. In this brief sketch we omit potential problems that may arise in such modeling, like obtaining disconnected networks or having 'blind spots', i.e., points in the polygon that cannot be cleared by using the above nodes and edges only. We say that a partial grid network G *covers* the polygon if G is connected and for each point p in the polygon there exist a node of G in distance at most r from p. See Figure 8 for an

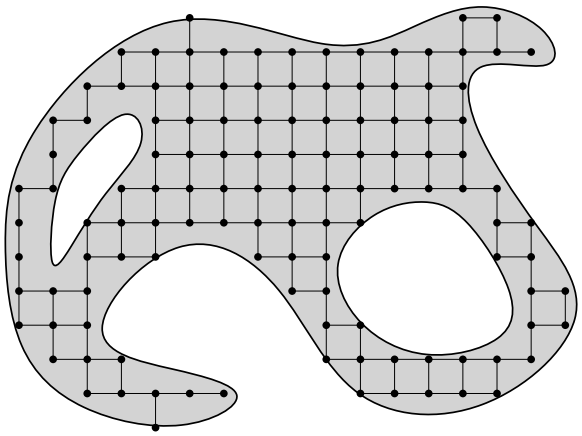

Figure 8: An example of the construction of a partial grid network.

example.

Note that any search strategy $\mathcal{S}'$ for a polygon P can be used to obtain a search strategy $\mathcal{S}$ for underlying partial grid network G as follows. For each searcher s used in $\mathcal{S}'$ introduce four searchers $s_1, \ldots, s_4$ that will 'mimic' its movements by going along edges of G. More precisely, the searchers $s_1, \ldots, s_4$ will ensure that at any point, if s is located at a point (x, y), then $s_1, \ldots, s_4$ will reside on nodes with coordinates $(\lfloor x/r \rfloor, \lfloor y/r \rfloor)$, $(\lfloor x/r \rfloor, \lceil y/r \rceil)$, $(\lceil x/r \rceil, \lfloor y/r \rfloor)$, $(\lceil x/r \rceil, \lceil y/r \rceil)$. In this way, at any time point, the four searchers in $\mathcal{S}$ protect an area that contains the area protected by s in $\mathcal{S}'$. This allows us to state the following.

Observation 7.1. *Let P be a polygon and let G by an underlying partial grid network that covers P. Then, there exists a search strategy for G using k searchers such that its execution in G results in clearing P and $k = O(p)$, where p is the minimum number of searchers required for clearing P (in continuous way).*

7.2 Open problems

In view of the lower bound shown in [23] that even in such simple networks as trees each distributed algorithm may be forced to use $\Omega(n/\log n)$ times more searchers than the connected search number of the underlying network, one possible line of research is to restrict attention to specific topologies that allow to obtain algorithms with good provable upper bounds. This work gives one such example. An interesting research direction is to find other non-trivial settings in which distributed search can be conducted efficiently.

The above questions related to network topologies can be stated more generally: what properties of the distributed model are crucial for such search for fast and invisible fugitive to be efficient? This work and also a recent one [5] suggest that a 'sense of direction' may be one such factor. Possibly interesting directions may be to analyze the influence of visibility on search scenarios.

We finally note that the only optimization criterion that was of interest in this work is the number of searchers. This coincides with the research done in offline search problems where this was the most important criterion giving nice ties between graph searching theory and structural graph theory. However, one may consider adding different optimization criteria like time (defined as the maximum number of synchronized steps) or cost (the total number of moves performed by all searchers).

An interesing research direction is to study agent guaranteed searching algorithms, in this or more general settings, under weaker assumptions regarding agent capabilities like memory size or ability to distinguish directions etc.

References

[1] L. Barrière, P. Flocchini, F.V. Fomin, P. Fraigniaud, N. Nisse, N. Santoro, and D.M. Thilikos. Connected graph searching. *Inf. Comput.*, 219:1–16, 2012.

[2] M.J. Best, A. Gupta, D.M. Thilikos, and D. Zoros. Contraction obstructions for connected graph searching. *Discrete Applied Mathematics*, DOI: 10.1016/j.dam.2015.07.036, 2015.

[3] L. Blin, J. Burman, and N. Nisse. Perpetual graph searching. Technical report, INRIA ⟨hal-00675233⟩, 2012.

[4] L. Blin, P. Fraigniaud, N. Nisse, and S. Vial. Distributed chasing of network intruders. *Theor. Comput. Sci.*, 399(1-2):12–37, 2008.

[5] P. Borowiecki, D. Dereniowski, and L. Kuszner. Distributed graph searching with a sense of direction. *Distributed Computing*, 28(3):155–170, 2015.

[6] J. Cai, P. Flocchini, and N. Santoro. Decontaminating a network from a black virus. *IJNC*, 4(1):151–173, 2014.

[7] T.H. Chung, G.A. Hollinger, and V. Isler. Search and pursuit-evasion in mobile robotics - A survey. *Auton. Robots*, 31(4):299–316, 2011.

[8] Y. Daadaa, P. Flocchini, and N. Zaguia. Network decontamination with temporal immunity by cellular automata. In *ACRI '10: Proc. of the 9th International Conference on Cellular Automata for Research and Industry, Ascoli Piceno, Italy, September 21-24, 2010*, pages 287–299, 2010.

[9] G. D'Angelo, A. Navarra, and N. Nisse. Gathering and exclusive searching on rings under minimal assumptions. In *ICDCN '14: Proc. of the 15th International Conference on Distributed Computing and Networking, Coimbatore, India, January 4-7, 2014*, pages 149–164, 2014.

[10] D. Dereniowski. Connected searching of weighted trees. *Theor. Comp. Sci.*, 412:5700–5713, 2011.

[11] D. Dereniowski. Approximate search strategies for weighted trees. *Theor. Comput. Sci.*, 463:96–113, 2012.

[12] D. Dereniowski. From pathwidth to connected pathwidth. *SIAM J. Discrete Math.*, 26(4):1709–1732, 2012.

[13] J.W. Durham, A. Franchi, and F. Bullo. Distributed pursuit-evasion without mapping or global localization via local frontiers. *Auton. Robots*, 32(1):81–95, 2012.

[14] J. Ellis and R. Warren. Lower bounds on the pathwidth of some grid-like graphs. *Discrete Applied Mathematics*, 156(5):545–555, 2008.

[15] P. Flocchini, M.J. Huang, and F.L. Luccio. Decontaminating chordal rings and tori using mobile agents. *Int. J. Found. Comput. Sci.*, 18(3):547–563, 2007.

[16] P. Flocchini, M.J. Huang, and F.L. Luccio. Decontamination of hypercubes by mobile agents. *Networks*, 52(3):167–178, 2008.

[17] P. Flocchini, F. Luccio, L. Pagli, and N. Santoro. Optimal network decontamination with threshold immunity. In *CIAC '13: Proc. of the 8th International Conference on Algorithms and Complexity, Barcelona, Spain, May 22-24, 2013*, pages 234–245, 2013.

[18] P. Flocchini, F. Luccio, L. Pagli, and N. Santoro. Network decontamination under m-immunity. *Discrete Applied Mathematics*, 201:114–129, 2016.

[19] F.V. Fomin and D.M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theor. Comput. Sci.*, 399(3):236–245, 2008.

[20] P. Fraigniaud, D. Ilcinkas, and A. Pelc. Oracle size: a new measure of difficulty for communication tasks. In *PODC'06: Proc. of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing*, pages 179–187, 2006.

[21] V.C.F. Gonçalves, P.M.V. Lima, N. Maculan, and F.M.G. França. A distributed dynamics for webgraph decontamination. In *ISoLA '10: Proc. of the 4th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation, Heraklion, Crete, Greece, October 18-21, 2010*, pages 462–472, 2010.

[22] G.A. Hollinger, S. Singh, J. Djugash, and A. Kehagias. Efficient multi-robot search for a moving target. *I. J. Robotic Res.*, 28(2):201–219, 2009.

[23] D. Ilcinkas, N. Nisse, and D. Soguet. The cost of monotonicity in distributed graph searching. *Distributed Computing*, 22(2):117–127, 2009.

[24] A. Kolling and S. Carpin. Multi-robot pursuit-evasion without maps. In *ICRA'10: Proc. of IEEE International Conference on Robotics and Automation*, pages 3045–3051, 2010.

[25] A.S. LaPaugh. Recontamination does not help to search a graph. *J. ACM*, 40(2):224–245, 1993.

[26] N. Megiddo, S.L. Hakimi, M.R. Garey, D.S. Johnson, and C.H. Papadimitriou. The complexity of searching a graph. *J. ACM*, 35(1):18–44, 1988.

[27] R. Moraveji, H. Sarbazi-Azad, and A.Y. Zomaya. Performance modeling of cartesian product networks. *J. Parallel Distrib. Comput.*, 71(1):105–113, 2011.

[28] N. Nisse and D. Soguet. Graph searching with advice. *Theor. Comput. Sci.*, 410(14):1307–1318, 2009.

[29] T.D. Parsons. Pursuit-evasion in a graph. In *Theory and Applications of Graphs, Lecture Notes in Mathematics*, volume 642, pages 426–441. Springer-Verlag, 1978.

[30] N.N. Petrov. A problem of pursuit in the absence of information on the pursued. *Differentsial'nye Uravneniya*, 18:1345–1352, 1982.

[31] E. Raboin, U. Kuter, and D.S. Nau. Generating strategies for multi-agent pursuit-evasion games in partially observable euclidean space. In *AAMAS '12: Proc. of the International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, June 4-8, 2012*, pages 1201–1202, 2012.

[32] C. Robin and S. Lacroix. Multi-robot target detection and tracking: taxonomy and survey. *Auton. Robots*, 40(4):729–760, 2016.

[33] S. Rodríguez, J. Denny, J. Burgos, A. Mahadevan, K. Manavi, L. Murray, A. Kodochygov, T. Zourntos, and N.M. Amato. Toward realistic pursuit-evasion using a roadmap-based approach. In *ICRA '11: Proc. of the IEEE International Conference on Robotics and Automation, Shanghai, China, 9-13 May 2011*, pages 1738–1745, 2011.

[34] S. Sachs, S.M. LaValle, and S. Rajko. Visibility-based pursuit-evasion in an unknown planar environment. *I. J. Robotic Res.*, 23(1):3–26, 2004.

[35] N.M. Stiffler and J.M. O'Kane. A complete algorithm for visibility-based pursuit-evasion with multiple pursuers. In *ICRA '14: Proc. of the IEEE International Conference on Robotics and Automation, Hong Kong, China, May 31 - June 7, 2014*, pages 1660–1667, 2014.

[36] B. Yang, D. Dyer, and B. Alspach. Sweeping graphs with large clique number. *Discrete Mathematics*, 309(18):5770–5780, 2009.